

## Performance Assessment of Computer Associates' CleverPath Portal



# Table of Contents

<b><i>Executive Summary .....</i></b>	<b><i>3</i></b>
<b><i>CleverPath Performance and Scalability Tests .....</i></b>	<b><i>4</i></b>
<i>Test Approach and Methodology .....</i>	<i>4</i>
<i>Test Environment .....</i>	<i>5</i>
<i>Test Process and Parameters Monitored.....</i>	<i>5</i>
<i>Test Results .....</i>	<i>6</i>
<i>Analysis of Test Results .....</i>	<i>13</i>
<b><i>Conclusion.....</i></b>	<b><i>15</i></b>
<b><i>Appendix: About Doculabs .....</i></b>	<b><i>15</i></b>

## EXECUTIVE SUMMARY

---

CleverPath Portal from Computer Associates (CA) is designed for organizations that need to unify their information sources and services via a common framework and interface. CA commissioned Doculabs, an independent analyst and consulting firm, to verify performance and scalability tests of the CleverPath solution. The benchmark testing was accomplished with a relatively small amount of hardware, using three Dell PowerEdge servers. The goal of this evaluation was to document CleverPath performance in several test scenarios based on predefined metrics of sustained throughput and response time.

Overall, Doculabs validated that CleverPath Portal provided good performance and scalability characteristics, making it a good fit for businesses with large or expanding user bases and associated service-level requirements. Doculabs validated that:

- **Adding servers improves throughput and response time.** One portal server serviced an average of 255.67 requests per second generated by 1,000 concurrent virtual user connections with an average response time of approximately 34.757 seconds. When a second server was added to the environment, the system handled an average of 533.61 requests per second with the same number of concurrent virtual user connections, and an average response time of approximately 16.252 seconds. When a third server was added, response time continued to improve, decreasing to approximately 10.088 seconds while handling an average of 840.95 requests per second with the same number of concurrent virtual user connections.
- **CleverPath Portal takes advantage of available hardware resources for scalability.** Adding more physical servers to the CleverPath Portal configuration yielded near linear scalability (when additional servers were added, the average number of user requests successfully serviced by the portal increased proportionally).
- **CPU utilization was nearly linear** for one, two, and three servers. The results showed that increasing number of physical servers resulted in a decrease in percentage of individual servers' CPU utilization and concurrently time to service requests.

# CLEVERPATH PERFORMANCE AND SCALABILITY TESTS

---

This section presents performance test results based on a recent test program validated by Doculabs. This section summarizes the test methodology, test results, and analysis.

## Test Approach and Methodology

The goal of this evaluation was to document CleverPath Portal's performance in several different test scenarios based on predefined metrics of sustained throughput and response time. The CA development group was seeking to the ability to service a peak load of 500 requests per second.

The only way to determine the limits of portal server performance and scalability is to run real-life tests that reflect the multitude of factors that can affect those attributes. Throughput can be impacted by any number of factors, including the number of concurrent user sessions, types of transactions (read vs. writes), available bandwidth, and network topology/configuration.

To determine if an enterprise portal solution is truly scalable when additional hardware resources are added to the environment, performance benchmarks must be conducted by actually configuring the hardware and/or software using a load balancer and switching device. Extrapolated results assume linear scalability, which is not necessarily the case in a real enterprise environment. Only the actual results obtained in this evaluation are included in this report.

The specific metrics that were measured in the testing were:

- ***Sustained Throughput*** – The maximum number of concurrent user requests serviced and maintained by the portal server. The number of actual requests serviced by the portal was determined from the statistics furnished by Microsoft Web Application Stress Tool.
- ***Response Time*** – The average time required to service user requests for composing dynamic web pages.

## Test Environment

The testing environment consisted of CleverPath Portal 3.5 running on one Dell PowerEdge and two Dell PowerEdge 6450 quad processor (700 megahertz) servers configured with 2 gigabytes of RAM. An Oracle 8.1.6 database ran on a domain within a Sun E1000 and had resources consisting of eight (400 megahertz) Sparc II processors with 8 gigabytes of RAM, a Cisco gigabit switch, and a Foundry Server Iron XL hardware-based load balancer and switching device. The Oracle database was accessed via the Cisco gigabit switch.

Client requests were generated using five Dell GX1 desktop machines. The load generating client machines, the Sun E10000 and Foundry Server Iron XL Load Balancer and switching device, were connected via the Cisco gigabit switch, and the portal servers were connected to the Foundry Server Iron XL Load Balancer and switching device. The following table shows the components used for testing.

Component	Configuration
Servers for CleverPath Portal 3.5 software	<ul style="list-style-type: none"><li>• One Dell PowerEdge 6400 Server (4 CPUs, 2 GB RAM) running Windows 2000 operating system</li><li>• Two Dell PowerEdge 6450 Servers (4 CPUs, 2 GB RAM) running Windows 2000 operating system</li><li>• One Sun E1000</li></ul>
Network Switch	<ul style="list-style-type: none"><li>• Cisco Catalyst 3500 Series XL</li></ul>
Hardware Load Balancer	<ul style="list-style-type: none"><li>• 1 Foundry Iron XL</li></ul>
Desktops	<ul style="list-style-type: none"><li>• Five Dell GX1</li></ul>
Database	<ul style="list-style-type: none"><li>• Oracle 8.1.6 database</li></ul>
User Database	<ul style="list-style-type: none"><li>• Configured with 1 million users and 100,000 user groups</li></ul>

Table 1 – Component Configuration

## Test Process and Parameters Monitored

The reference performance test utilized one Dell PowerEdge 6400 and two Dell PowerEdge 6450 quad processor servers – the basic test configuration to determine if the CleverPath Portal solution could service 500 concurrent user requests generated by five Dell GX1 desktop clients. Doculabs monitored the following parameters during the performance assessment:

- **Throughput** – The number of concurrent requests per second being serviced
- **Total system CPU utilization** – The amount of work that the CPU had to perform in order to service and maintain users based on total user load.
- **Response time** – The time required for the client to receive the last byte of information requested.

## Test Results

This section summarizes the results of the performance tests for each of the areas and scenarios evaluated. The dependent variables measured in all cases are throughput and CPU utilization. For performance tests, temporal profiles of these parameters were examined. To investigate scalability, these variables were examined in terms of their behavior as a function of a number of physical servers.

Number of Servers	Concurrent User Connections	Average Throughput Requests/Sec	Response Time (Seconds)	CPU Utilization (%)
3	10	78.24	00.074	03.35
3	50	270.94	00.115	13.98
3	100	486.19	00.112	25.12
3	250	1033.83	00.117	66.52
3	500	1207.04	00.340	84.58
3	750	1211.06	00.554	65.40
3	1000	1214.63	00.756	87.46

Table 2 - CleverPath Portal Throughput Performance – Baseline

Number of Servers	Concurrent User Connections	Average Throughput Requests/Sec	Response Time (Seconds)	CPU Utilization (%)
1	10	92.26	00.348	20.19
1	50	266.72	01.269	84.98
1	100	257.27	03.091	85.49
1	250	254.70	06.111	86.75
1	500	254.28	17.277	87.14
1	750	254.32	26.044	85.69
1	1000	255.67	34.757	87.08

Table 3 – CleverPath Portal Throughput Performance - One Server

Number of Servers	Concurrent User Connections	Average Throughput Requests/Sec	Response Time (Seconds)	CPU Utilization (%)
2	10	102.37	00.339	10.43
2	50	346.02	00.655	41.20
2	100	538.13	01.166	76.98
2	250	542.40	03.647	86.74
2	500	543.96	07.757	85.93
2	750	544.83	11.831	84.88
2	1000	533.61	16.352	85.87

*Table 4 – CleverPath Portal Throughput Performance - Two Servers*

Number of Servers	Concurrent User Connections	Average Throughput Requests/Sec	Response Time (Seconds)	CPU Utilization (%)
3	10	99.53	00.306	01.79
3	50	335.02	00.619	20.35
3	100	613.87	00.893	41.28
3	250	857.26	02.119	66.17
3	500	842.83	04.757	64.99
3	750	839.55	07.450	72.58
3	1000	840.95	10.088	67.07

*Table 5 - CleverPath Portal Throughput Performance - Three Servers*

As the tables show, CleverPath Portal's throughput and CPU utilization increased in a near-linear manner as the number of servers increased to three servers. Throughput and CPU utilization increased proportionately as servers were gradually added to the configuration.

The following figure shows the average number of requests per second serviced by a three-server configuration, with a varying number of concurrent user connections accessing a baseline configuration with an empty workplace.




Figure 1 - Average Throughput for Baseline Scenario

The following figure shows the average number of requests per second serviced by different numbers of servers, with a varying number of concurrent user connections accessing a populated workplace.




Figure 2 - Scalability and Performance



The following figure shows the average response time for different numbers of servers with a varying number of concurrent user connections.




Figure 3 - Response Time

The following figure shows the temporal relationship between throughput and CPU utilization with one server.




Figure 4 - Throughput and CPU Utilization – One Server

The following figure shows the temporal relationship between throughput and CPU utilization with two servers.




Figure 5 - Throughput and CPU Utilization – Two Servers

The following figure shows the temporal relationship between throughput and CPU utilization with three servers.




Figure 6 - Throughput and CPU Utilization – Three Servers

The following figure shows the temporal relationship between response time and CPU utilization with one server.




Figure 7 - Response Time and CPU Utilization – One Server

The following figure shows the temporal relationship between response time and CPU utilization with two servers.




Figure 8 - Response Time and CPU Utilization – Two Servers

The following figure shows the temporal relationship between response time and CPU utilization with three servers.




Figure 9 - Response Time and CPU Utilization – Three Servers

## Analysis of Test Results

This section contains Doculabs' analysis of the test results in terms of performance and throughput, and includes a scalability discussion on each of these parameters.

Scalability is essential for enterprise-class products, especially in terms of their ability to support an organization's future growth – a scaleable system will be able to support increases in user workloads without sacrificing functionality or performance. However, designing a scalable portal infrastructure requires an estimate of potential users, future users, and information and application sources, as well as an understanding of the system resource growth required to support the environment.

### *Performance Analysis*

As indicated in Table 3, one portal server serviced an average of 255.67 requests per second generated by 1,000 concurrent virtual user connections with an average response time of approximately 34.757 seconds. Table 4 demonstrates that when a second server was added to the environment, the average response time dropped to approximately 16.252 seconds, with the same number of concurrent virtual user connections. When a third server was added, response time continued to improve, decreasing to approximately 10.088 seconds with the same number of concurrent virtual user connections, as illustrated in Table 5.

Figures 7 through 9 illustrate the output from the Microsoft Web Application Stress Tool, which demonstrates that response time decreased proportionately as physical servers were added to the portal environment. The average CPU utilization for 1,000 concurrent user connections also consistently decreased as servers were added. With 3 servers, CPU utilization was approximately 67.07 percent throughout the test, and never rose above 72.58 percent.

As shown in Figures 4 through 6, CPU utilization was nearly linear for one, two, and three servers. An inspection of the output of the Windows Performance Monitor utility, which monitors per-processor activity during user activity, demonstrates that CleverPath Portal takes advantage of multi-server configurations. This means that an increase in the number of physical servers will decrease the percentage of CPU utilization of individual servers and concurrently decrease the time to service requests; thus, servers will be available to service other subsequent requests.

However, it is important to note that in a production environment, portal performance can be impacted by any number of factors, including the number of concurrent user sessions, types of transactions (read vs. writes), available bandwidth, and network topology/configuration. The test suite made serial requests of the portal server in order to retrieve portlet information that was contained in separate iFrames. Doculabs expects that this resulted in degraded response times, as a typical production environment would process requests in parallel using multi-threading to achieve the appearance of concurrency. Conversely, the test suite did not make extensive use of back office integrations with products such as ERP or CRM systems, which would have introduced an additional level of complexity and variability to the benchmarking process.

### ***Throughput Analysis***

In the first test, using one server, CleverPath Portal successfully serviced an average of 255.67 requests per second with a 1,000-user workload. In the second test an average of 533.61 requests per second with a 1,000-user workload was successfully serviced utilizing two servers. Likewise, utilizing three servers, an average of 840.95 requests per second was serviced in the second test with a 1,000-user workload.

As the data shows, adding more physical servers to the portal configuration yields linear scalability. When additional servers were added, the average number of user requests successfully serviced by the portal increased. Both CPU utilization and response time decreased proportionately with the number of servers. In addition, the average number of requests per second being serviced by the portal increased proportionately to the number of portal servers available to handle requests.

Overall, the benchmark results prove that CleverPath Portal takes advantage of available hardware resources to achieve near linear scalability, and clearly illustrate a scalability factor of approximately one.

## CONCLUSION

---

CleverPath Portal is built on a strong Java-based foundation and core group of services that allow it to effectively meet the needs of many organizations as they begin to roll out their portal initiatives. As enterprise needs evolve, organizations can take advantage of the wide variety of complementary Computer Associates products designed to extend CleverPath Portal's capabilities.

Overall, CleverPath Portal's underlying product architecture is robust, and has demonstrated its ability to scale in a linear fashion as additional hardware is added, when used with an industry standard application server such as IBM WebSphere. This should be an attractive characteristic for operations managers that anticipate expanding user bases and associated service-level requirements.

## APPENDIX: ABOUT DOCULABS

---

Doculabs, Inc. is an independent research and consulting firm that improves the way companies plan for, select and optimize emerging technologies through project-based services. Based in Chicago, Doculabs' consulting services are grounded in research that combines hands-on evaluation of technology with real-time business knowledge gained from engagements with Fortune 1000 clients. Doculabs' services help clients deliver on their business strategies through solutions in areas such as enterprise content management, relationship management and infrastructure. Doculabs' consulting services are completely objective because the firm does not sell software or integration services. Clients benefit from more informed decisions in a shorter period of time, which reduces risk, lowers costs, and improves time-to-market.

For more information about Doculabs, visit the Web site at [www.doculabs.com](http://www.doculabs.com) or call (312) 433-7793.



**1201 West Harrison Street, Third Floor**  
**Chicago, IL 60607**  
**312-433-7793**  
**[www.doculabs.com](http://www.doculabs.com)**